



2012 Programming weather, climate, and earth-system models on heterogeneous multi-core platforms

September 12-13, 2012 at the National Center for Atmospheric Research in Boulder, Colorado

Addressing the Increasing Challenges of Debugging on Accelerated Multi-Core Systems

Ed Hinkel
Senior Sales Engineer

Agenda

Overview - Rogue Wave & TotalView
Heterogeneous Systems - Then and Now
Debugging Accelerated Systems
GPU - Nvidia CUDA
MIC - Intel Phi

Rogue Wave Today



The largest independent provider of cross-platform software development tools and embedded components for the next generation of HPC applications

Visual Numerics®

Leader in embeddable math and statistics algorithms and visualization software for data-intensive applications.



Leading provider of intelligent software technology which analyzes and optimizes computing performance in single and multi-core environments.



Industry-leading interactive analysis and debugging tools for the world's most sophisticated software applications.

Latest addition to the Rogue Wave family: **Rogue Wave Visualization for C++**
(Formerly IBM's ILOG Visualization for C++ products)

Representative Customers




Heterogeneous Systems and the Need for Speed

A couple of pioneers

**Remember
SiCortex?**

27-Node Module



14

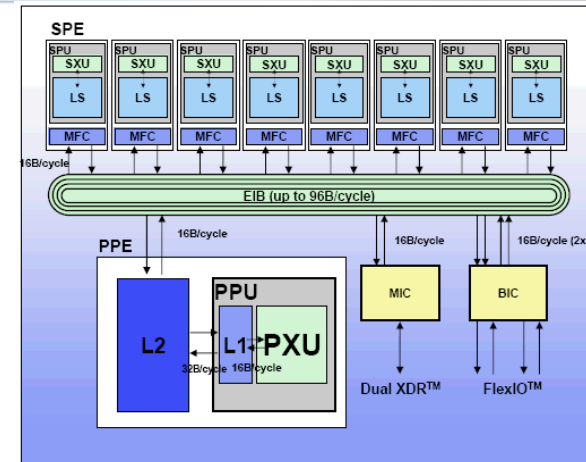
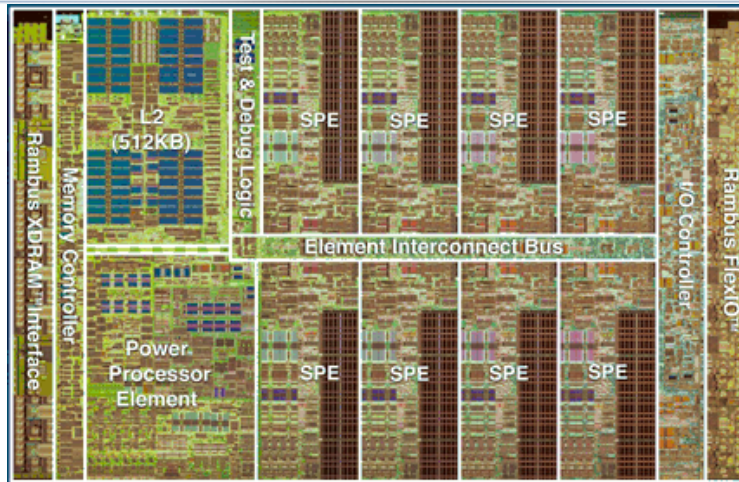


An X-86 host with up to 972 MIPS nodes, with up to 5,832 cores and 7,776 GB of memory

**Remember
Cell?**


Heterogeneous Systems - CELL

March 2001 - Nov 2009



Source: M. Gochwind et al., Hot Chips-17, August 2005

RoadRunner

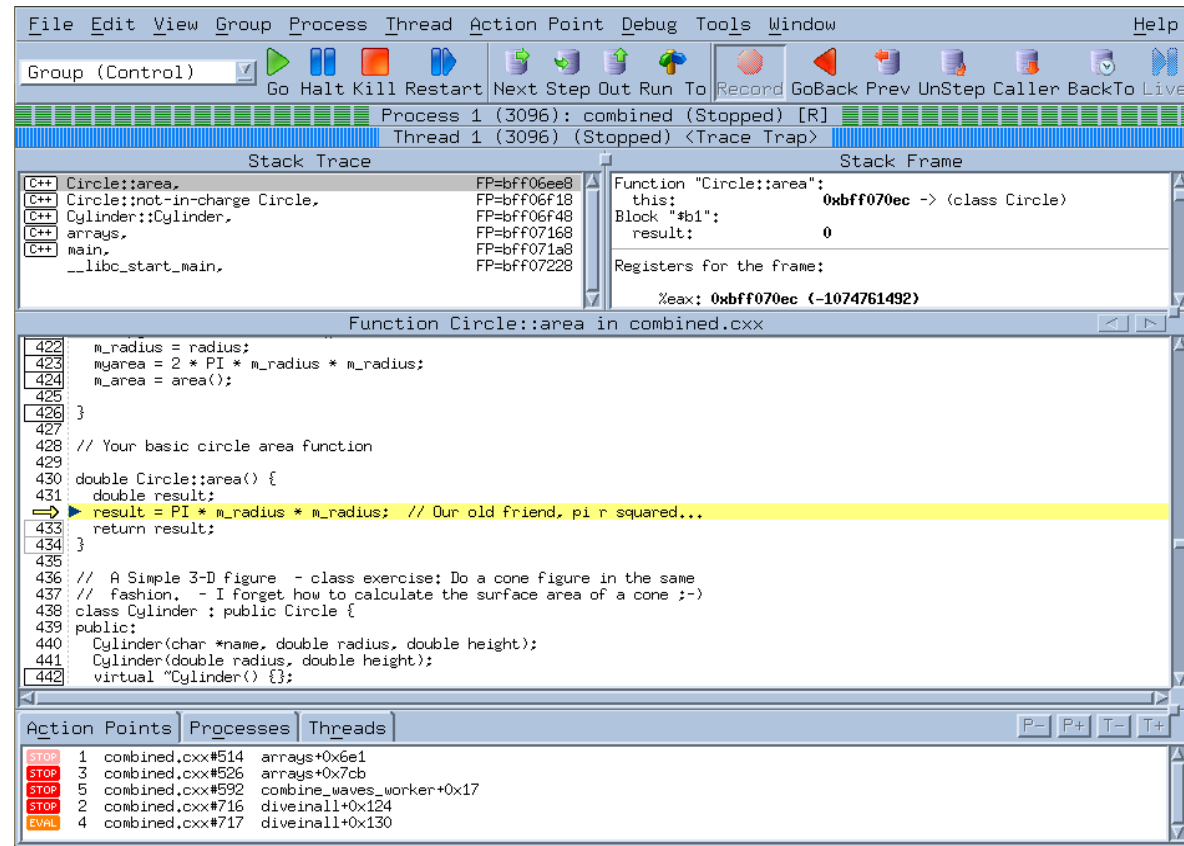


a hybrid design with 12,960 [IBM PowerXCell 8i](#) and 6,480 [AMD Opteron](#) dual-core processors

What is TotalView?

A comprehensive debugging solution for demanding parallel, heterogeneous, and multi-core applications

- **Wide compiler & platform support**
 - C, C++, Fortran 77 & 90, UPC
 - Unix, Linux, OS X
- **Handles Concurrency**
 - Multi-threaded Debugging
 - Multi-process Debugging
- **Integrated Memory Debugging**
- **Reverse Debugging for Linux**
- **Supports Multiple Usage Models**
 - Powerful and Easy GUI – Highly Graphical
 - CLI for Scripting
 - Long Distance Remote Debugging
 - Unattended Batch Debugging

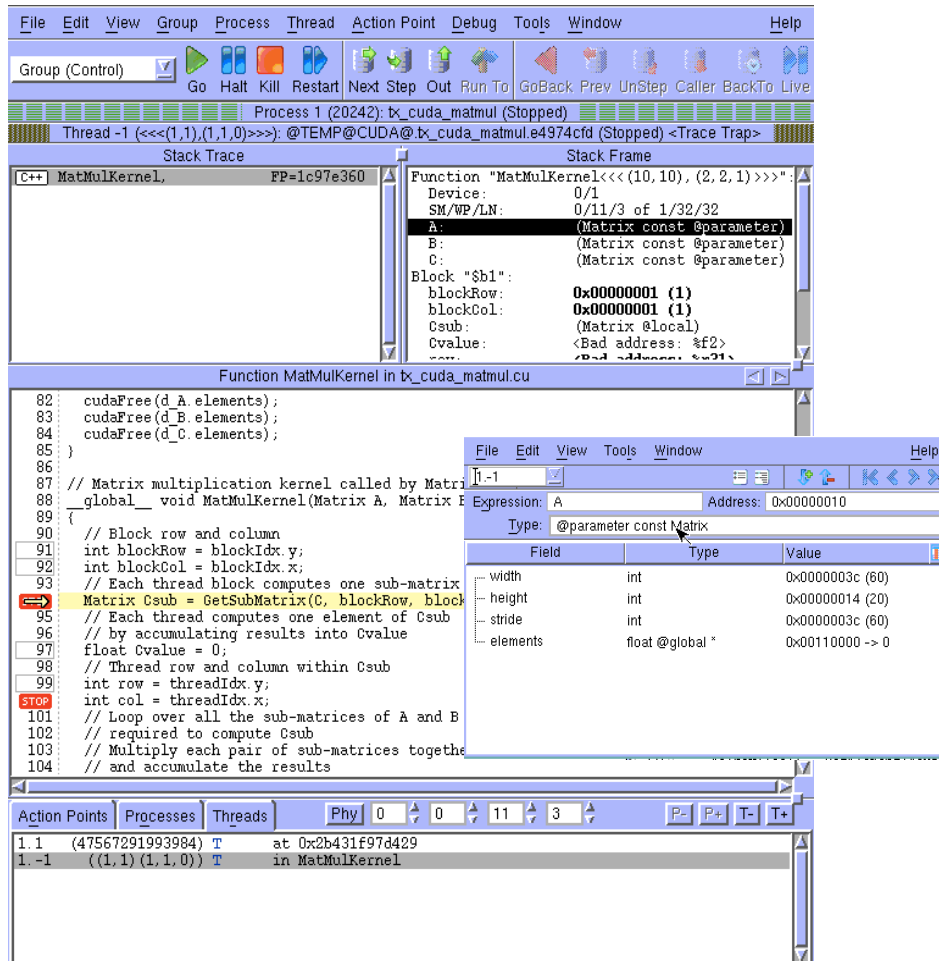


GPU Debugging

with

TotalView

CUDA Port of TotalView



Full visibility of both Linux and GPU threads

Device threads shown as part of the parent Unix process

Handles all the differences between the CPU and GPU

Fully represent the hierarchical memory

Display data at any level (registers, local, block, global or host memory)

Making it clear where data resides with type qualification

Thread and Block Coordinates

Built in runtime variables display threads in a warp, block and thread dimensions and indexes

Displayed on the interface in the status bar, thread tab and stack frame

Device thread control

Warps advance synchronously

Handles CUDA function inlining

Step into or over inlined functions

Functions show on stack trace

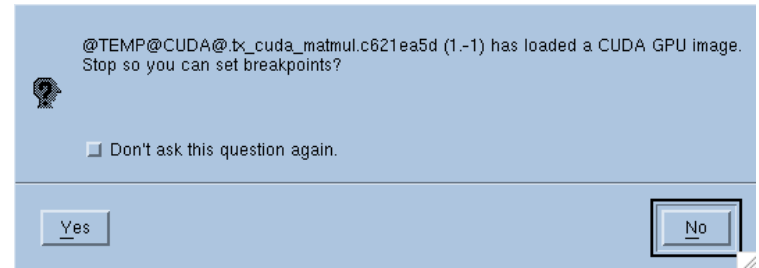
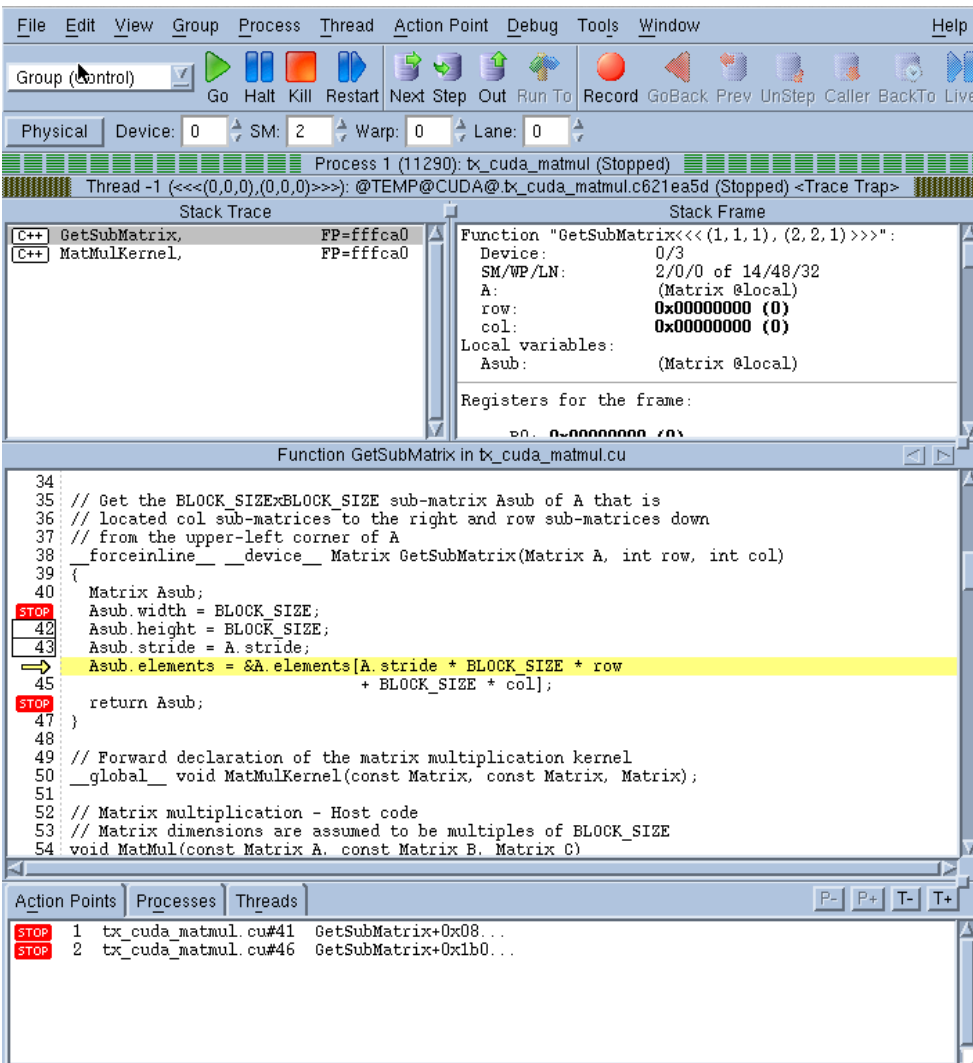
Reports memory access errors

CUDA memcheck

Multi-Device Support

Can be used with MPI

Starting TotalView



- Once a new kernel is loaded TotalView provides the option to stop and set breakpoints
- TotalView automatically configures the GUI for CUDA debugging
- Debugging CUDA code is done by using normal TotalView commands and procedures

GPU Device Status Display

Provides the “high-level” view

- Values automatically update as you step through code
- Shows what hardware is in use
- Helps to map between logical and hardware coordinates

Name	Description
Device 0/3	
Device Type	gf100
Lanes	32
SM 2/1	
Valid Warps	0000000000000001
Warp 00/48	Block (0,0,0)
Lane 00/32	Thread (0,0,0)
LPC	0000000019aa94d8
Lane 01/32	Thread (1,0,0)
LPC	0000000019aa94d8
Lane 02/32	Thread (2,0,0)
LPC	0000000019aa94f0
Lane 03/32	Thread (3,0,0)
LPC	0000000019aa94f0
Lane 04/32	Thread (4,0,0)
LPC	0000000019aa94f0
Lane 05/32	Thread (5,0,0)
LPC	0000000019aa94f0
Lane 06/32	Thread (6,0,0)
LPC	0000000019aa94f0
Lane 07/32	Thread (7,0,0)
LPC	0000000019aa94f0
Lane 08/32	Thread (8,0,0)
LPC	0000000019aa94f0
Lane 09/32	Thread (9,0,0)
LPC	0000000019aa94f0
Valid/Active/Divergent	000003ff, 000003fc, 00000003
SM Type	sm_20
SMs	14
Warps	48
Device 1/3	
Device Type	gt200
Lanes	32
SM Type	sm_13

GPU Device Status Display

Provides detailed information for:

Device and Type

SMs

Warps

Lanes with PC

Name	Description
Device 0/3	
Device Type	gf100
Lanes	32
SM 2/1	
Valid Warps	000000
Warp 00/48	Block
Lane 00/32	Thread
LPC	000000
Lane 01/32	Thread
LPC	000000
Lane 02/32	Thread
LPC	000000
Lane 03/32	Thread
LPC	000000
Lane 04/32	Thread
LPC	000000
Lane 05/32	Thread

Information updates as you step

GPU Device Status Display

Name	Description
Device 0/3	
Device Type	gf100
Lanes	32
SM 2/1	
Valid Warps	0000000000000001
Warp 00/48	Block (0,0,0)
Lane 00/32	Thread (0,0,0)
LPC	0000000019aa94d8
Lane 01/32	Thread (1,0,0)
LPC	0000000019aa94d8
Lane 02/32	Thread (2,0,0)
LPC	0000000019aa94f0
Lane 03/32	Thread (3,0,0)
LPC	0000000019aa94f0
Lane 04/32	Thread (4,0,0)
LPC	0000000019aa94f0
Lane 05/32	Thread (5,0,0)
LPC	0000000019aa94f0
Lane 06/32	Thread (6,0,0)
LPC	0000000019aa94f0
Lane 07/32	Thread (7,0,0)
LPC	0000000019aa94f0
Lane 08/32	Thread (8,0,0)
LPC	0000000019aa94f0
Lane 09/32	Thread (9,0,0)
LPC	0000000019aa94f0
Valid/Active/Divergent	000003ff, 000003fc, 00000003
SM Type	sm_20
SMs	14
Warps	48
Device 1/3	
Device Type	gt200
Lanes	32
SM Type	sm_13

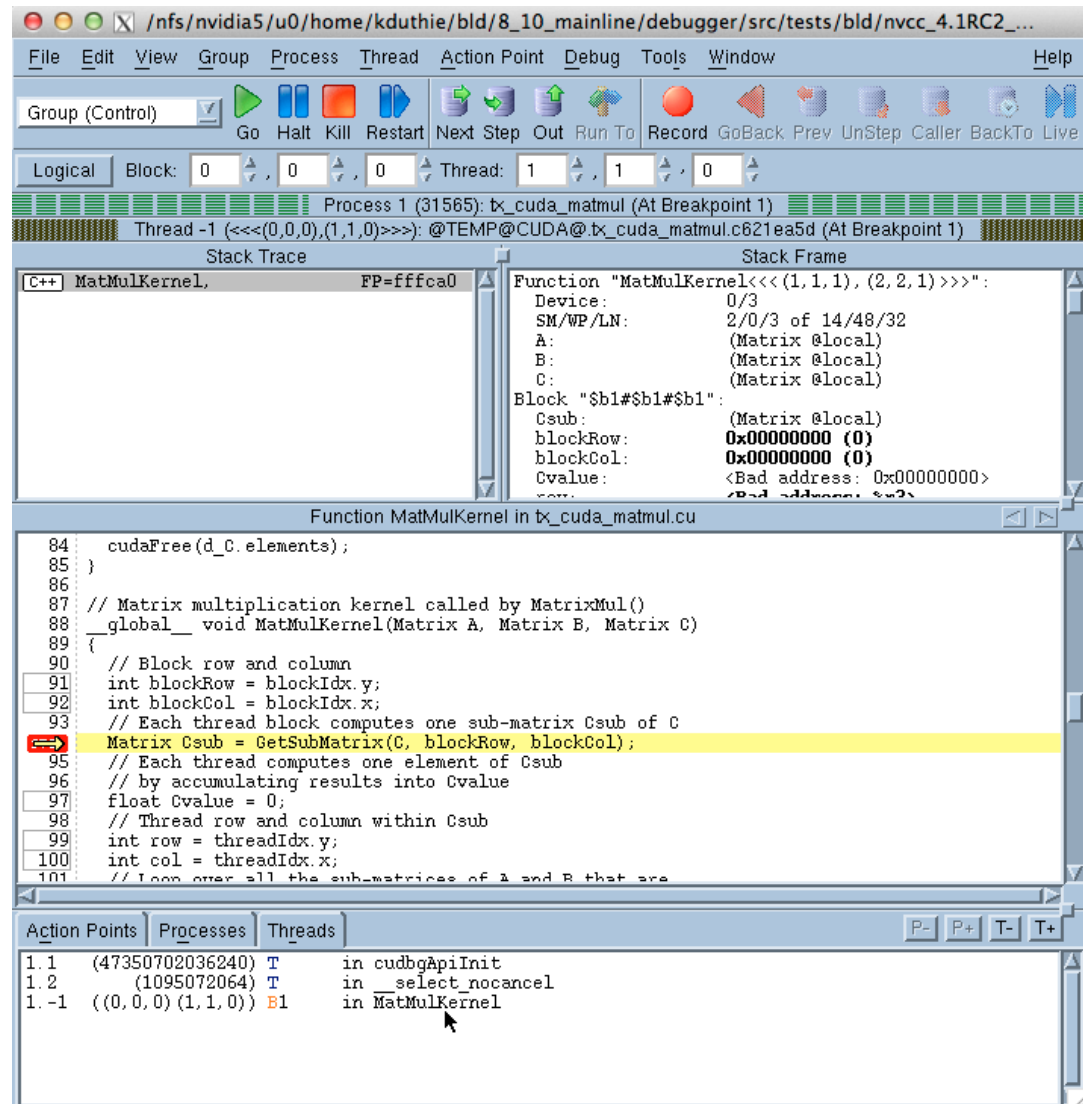
It also provides information
for divergent GPU threads

Different PC for two
groups of Lanes

State of Lanes
inside the Warp

Debugging CUDA

Information on
GPU execution,
location and
data is readily
available.
... the same as it
is for Linux
processes and
threads.



Debugging CUDA

e/kduthie/bld/8_10_mainline/debugger/src/tests/bld/nvcc_4.1RC2_...

Thread Action Point Debug Tools Window Help

Restart Next Step Out Run To Record GoBack Prev UnStep Caller BackTo Live

0 Thread: 1 1 0

Process 1 (31565): tx_cuda_matmul (At Breakpoint 1)

1,0)>>>: @TEMP@CUDA@tx_cuda_matmul.c621ea5d (At Breakpoint 1)

Stack Frame

FP=fffc0

Function "MatMulKernel<<<(1,1,1), (2,2,1)>>>":

Device: 0/3

SM/WP/LN: 2/0/3 of 14/48/32

A: (Matrix @local)

B: (Matrix @local)

C: (Matrix @local)

Block "\$b1#\$b1#\$b1":

Csub: (Matrix @local)

blockRow: 0x00000000 (0)

blockCol: 0x00000000 (0)

Cvalue: <Bad address: 0x00000000>

... <Bad address: 8x2>

ction MatMulKernel in tx_cuda_matmul.cu

kernel called by MatrixMul()

nel(Matrix A, Matrix B, Matrix C)

y;

X;

computes one sub-matrix Csub of C

rix(C, blockRow, blockCol);

one element of Csub

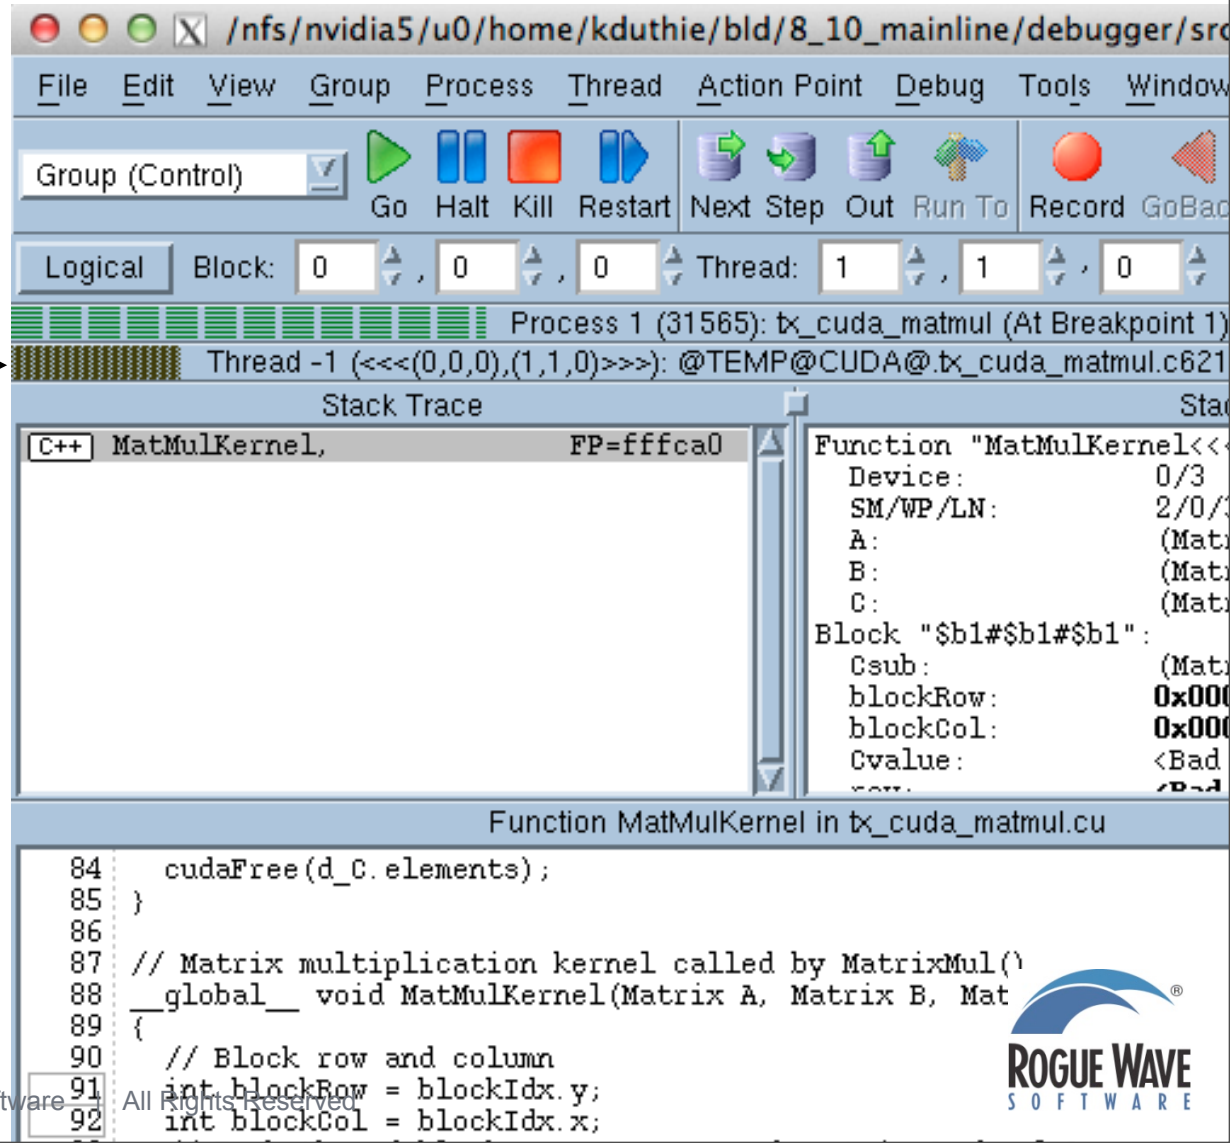
lts into Cvalue

CUDA grid and block dimensions, lanes/warp, warps/SM,

Parameter, register, local and shared variables

Debugging CUDA

GPU focus
thread logical
coordinates in
the header...



The screenshot shows the Rogue Wave debugger interface. The top menu bar includes File, Edit, View, Group, Process, Thread, Action Point, Debug, Tools, and Window. Below the menu is a toolbar with buttons for Go, Halt, Kill, Restart, Next Step, Out, Run To, Record, and Go Back. The 'Logical' tab is selected, showing the thread coordinates as Block: 0, 0, 0 and Thread: 1, 1, 0. The 'Stack Trace' pane shows the function 'MatMulKernel' with FP=fffca0. The 'Function MatMulKernel in tx_cuda_matmul.cu' pane shows the source code for the kernel, with the current line of execution highlighted at line 92.

Process 1 (31565): tx_cuda_matmul (At Breakpoint 1)

Thread -1 (<<<(0,0,0),(1,1,0)>>>): @TEMP@CUDA@tx_cuda_matmul.c621

Stack Trace

Function	FP
MatMulKernel	fffca0

Function MatMulKernel in tx_cuda_matmul.cu

```
84  cudaFree(d_C.elements);
85  }
86
87  // Matrix multiplication kernel called by MatrixMul()
88  __global__ void MatMulKernel(Matrix A, Matrix B, Matrix C)
89  {
90      // Block row and column
91      int blockDim = blockDim.y;
92      int blockDim = blockDim.x;
```

Debugging CUDA

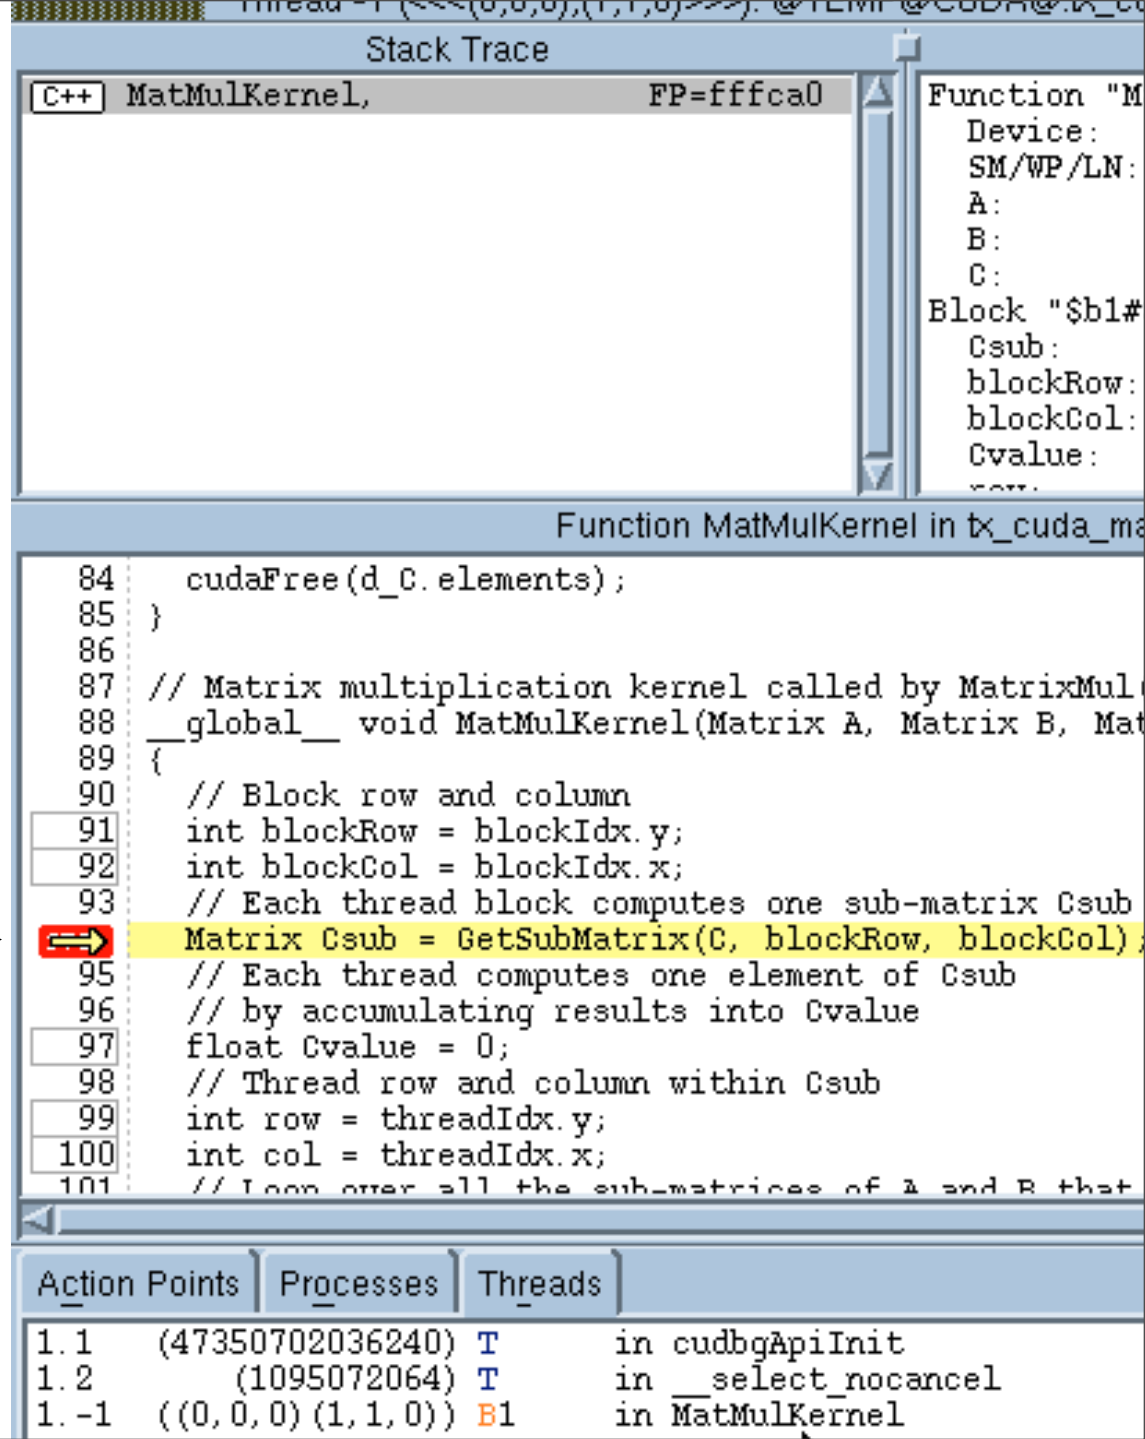
... as well as in
the Process
Window

```
87 // Matrix multiplication kernel called by Matr
88 __global__ void MatMulKernel(Matrix A, Matrix
89 {
90     // Block row and column
91     int blockRow = blockIdx.y;
92     int blockCol = blockIdx.x;
93     // Each thread block computes one sub-matrix
94     Matrix Csub = GetSubMatrix(C, blockRow, block
95     // Each thread computes one element of Csub
96     // by accumulating results into Cvalue
97     float Cvalue = 0;
98     // Thread row and column within Csub
99     int row = threadIdx.y;
100    int col = threadIdx.x;
101    // Loop over all the sub-matrices of A and B
```

Action Points	Processes	Threads
1.1	(47350702036240)	T in cudbgApiInit
1.2	(1095072064)	T in __select_nocancel
1.-1	((0, 0, 0) (1, 1, 0))	B1 in MatMulKernel

Debugging CUDA

PC arrow shows the
Program Counter
for the warp



The screenshot displays the debugger's interface for the `MatMulKernel` function. The **Stack Trace** pane at the top shows the current frame: `C++ MatMulKernel, FP=fffca0`. The **Function MatMulKernel in tx_cuda_ma** pane shows the source code with line numbers 84 to 101. Line 93, `Matrix Csub = GetSubMatrix(C, blockRow, blockCol);`, is highlighted in yellow, and a red arrow points to it from the callout box. The **Action Points** pane at the bottom shows the execution flow:

Action Points	Processes	Threads
1. 1	(47350702036240)	T in cudbgApiInit
1. 2	(1095072064)	T in __select_nocancel
1. -1	((0, 0, 0) (1, 1, 0))	OB1 in MatMulKernel

Debugging CUDA

The screenshot displays the Rogue Wave debugger interface. At the top, the file path is `/nfs/nvidia5/u0/home/kduthie/bld/8_10_mainline/debugg`. The menu bar includes `File`, `Edit`, `View`, `Group`, `Process`, `Thread`, `Action Point`, `Debug`, and `Tools`. The toolbar contains buttons for `Go`, `Halt`, `Kill`, `Restart`, `Next Step`, `Out`, `Run To`, and `Record`. Below the toolbar, the `Logical` tab is selected, showing `Block: 0, 0, 0` and `Thread: 1, 1, 0`. The `Stack Trace` window shows the `MatMulKernel` function at `FP=fffca0`. A callout box for the function shows `Device: SM/WP/LN: A: B: C: Block "$b1#$b1#$b1"`. The `Variable Window` is open, showing the variable `A` at `Address: 0x00000010` with type `@parameter const Matrix`. The variable `A` is expanded, showing fields `width`, `height`, `stride`, and `elements` with their respective values. The `elements` field has a value of `0x00110000 -> 0`. The `Source` window shows the `MatMulKernel` function code, with line 94 highlighted. A red arrow points from the `Variable Window` to the `Source` window.

Group (Control) [v] Go Halt Kill Restart Next Step Out Run To Record

Logical Block: 0, 0, 0 Thread: 1, 1, 0

Process 1 (31565): tx_cuda_matmul (At Break)

Thread -1 (<<<(0,0,0),(1,1,0)>>>): @TEMP@CUDA@tx_cuda_matmul

Stack Trace

C++ MatMulKernel, FP=fffca0

Function "MatMulKernel"

Device:

SM/WP/LN:

A:

B:

C:

Block "\$b1#\$b1#\$b1"

A - MatMulKernel - 1.-1

File Edit View Tools Window Help

1.-1

Expression: A Address: 0x00000010

Type: @parameter const Matrix

Field	Type	Value
width	int	0x00000002 (2)
height	int	0x00000002 (2)
stride	int	0x00000002 (2)
elements	float @global *	0x00110000 -> 0

84
85
86
87 //
88 g
89 {
90 /
91 ir
92 ir
93 /
94 M
95 /
96 /
97 f
98 /
99 ir
100 ir
101 // Loop over all the sub-matrices of A and B that are

Dive on any variable name to open a variable window

Debugging CUDA

■ A - MatMulKernel - 1.-1

File Edit View Tools Window Help

1.-1

Expression: A Address: 0x00000010

Type: @parameter const Matrix

Field	Type	Value
width	int	0x00000002 (2)
height	int	0x00000002 (2)
stride	int	0x00000002 (2)
elements	float @global *	0x00110000 -> 0

“@parameter” type qualifier indicates that variable “A” is in parameter storage

Address 0x10 is an offset within parameter storage

Pointer value 0x110000 is an offset within global storage

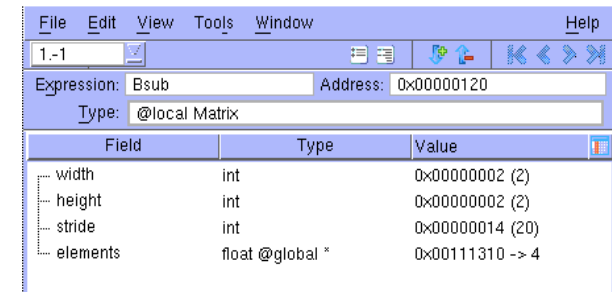
“elements” is a pointer to a float in global storage

Storage Qualifiers

- **Denotes location in hierarchical memory**
 - Part of the type – using “@” notation
 - Each memory space has a separate address space so 0x00001234 could refer to several places

Storage Qualifier	Meaning
@parameter	Address is an offset within parameter storage.
@local	Address is an offset within local storage.
@shared	Address is an offset within shared storage.
@constant	Address is an offset within constant storage.
@global	Address is an offset within global storage.
@register	Address is a PTX register name

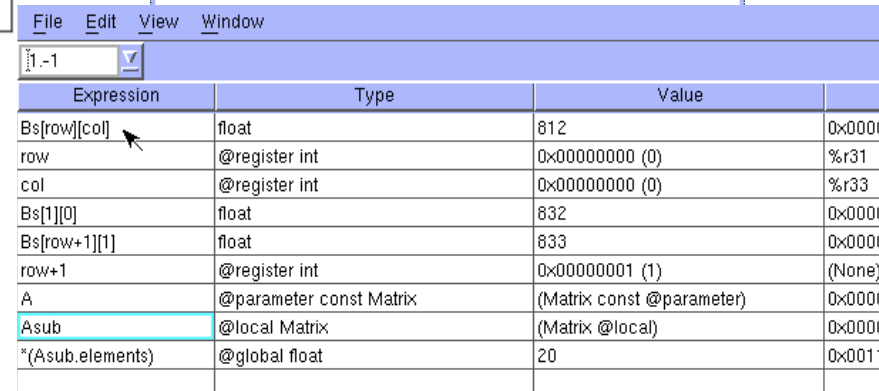
- **Used throughout expression system**
 - You can cast to switch between different spaces



The screenshot shows a code editor window with a menu bar (File, Edit, View, Tools, Window, Help) and a toolbar. Below the toolbar, there's a text area with the following content:

```
1.-1  
Expression: Bsub Address: 0x00000120  
Type: @local Matrix
```

Field	Type	Value
width	int	0x00000002 (2)
height	int	0x00000002 (2)
stride	int	0x00000014 (20)
elements	float @global *	0x00111310 -> 4



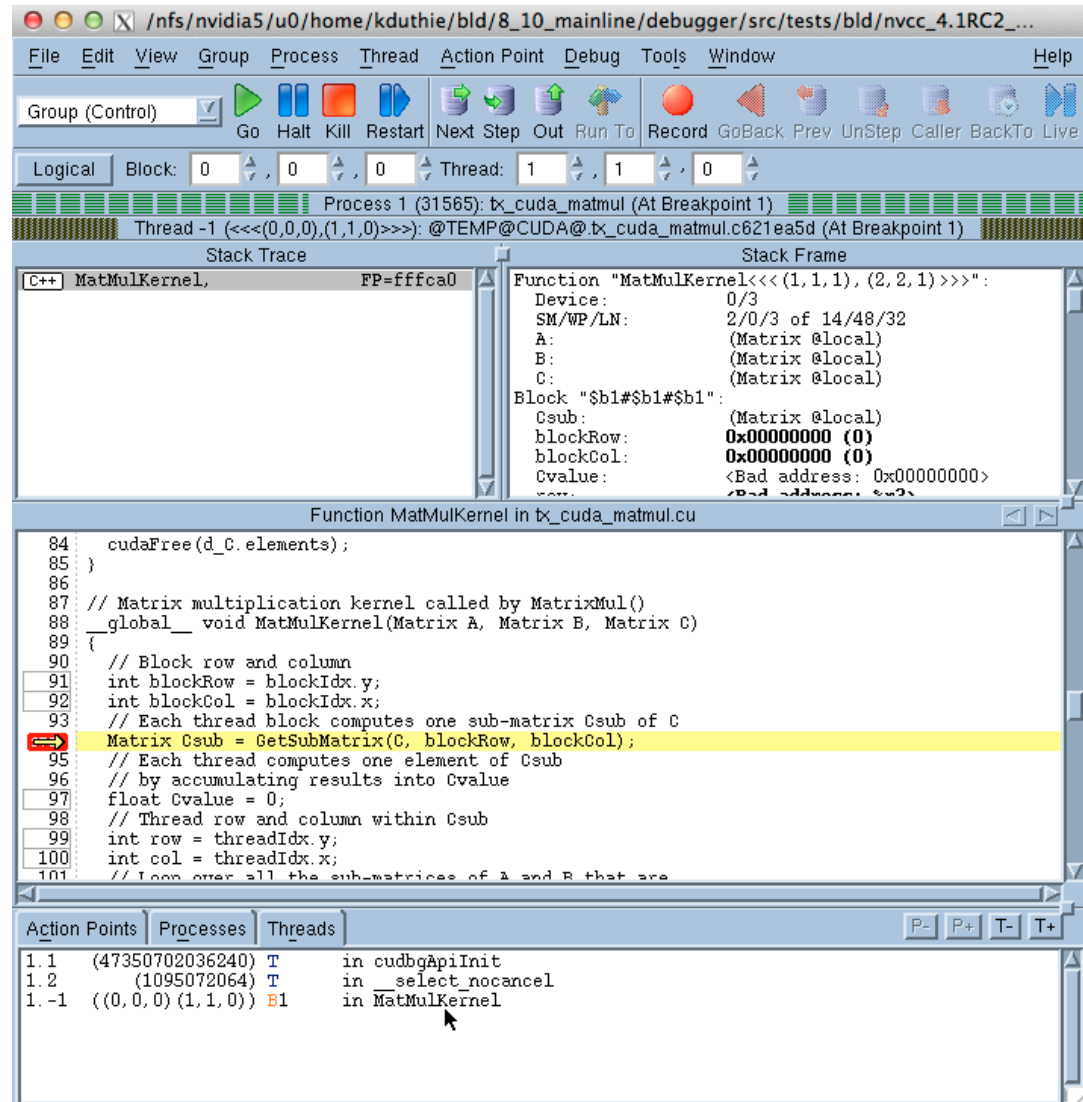
The screenshot shows a code editor window with a menu bar (File, Edit, View, Window) and a toolbar. Below the toolbar, there's a text area with the following content:

```
1.-1  
Expression: Bsub Address: 0x00000120  
Type: @local Matrix
```

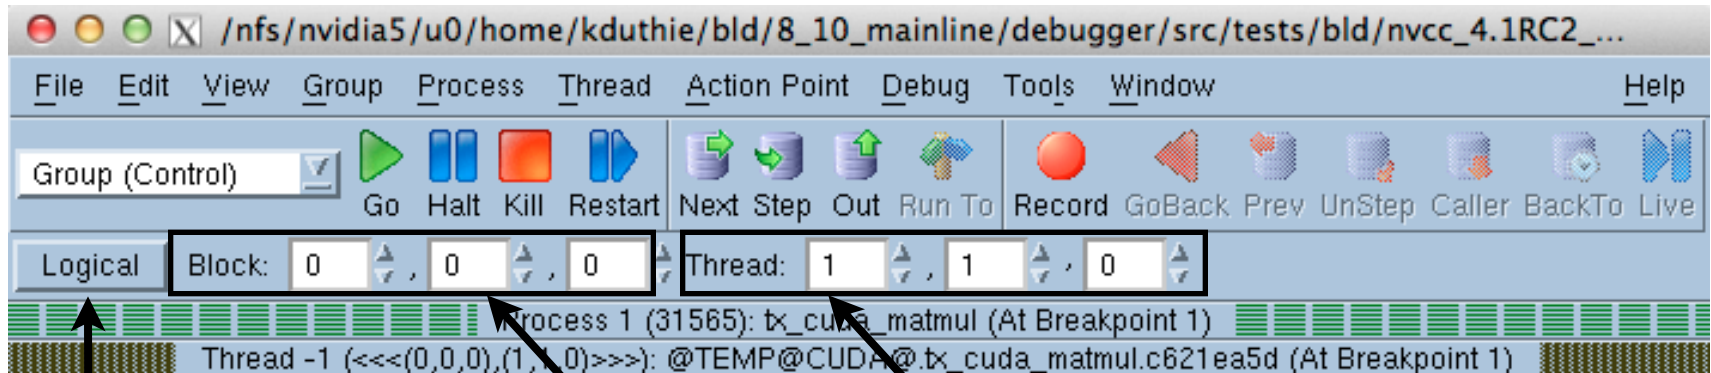
Expression	Type	Value	
Bs[row][col]	float	812	0x0000000000000000
row	@register int	0x00000000 (0)	%r31
col	@register int	0x00000000 (0)	%r33
Bs[1][0]	float	832	0x0000000000000000
Bs[row+1][1]	float	833	0x0000000000000000
row+1	@register int	0x00000001 (1)	(None)
A	@parameter const Matrix	(Matrix const @parameter)	0x0000000000000000
Asub	@local Matrix	(Matrix @local)	0x0000000000000000
*(Asub.elements)	@global float	20	0x0000000000000000

Debugging CUDA - Navigation

Navigate
through your
CUDA code in
the Process
Window as you
wish...
Using either of
two coordinate
systems:



Debugging CUDA - Navigation



CUDA GPU threads have a negative TotalView thread ID

Block
(x,y,z)

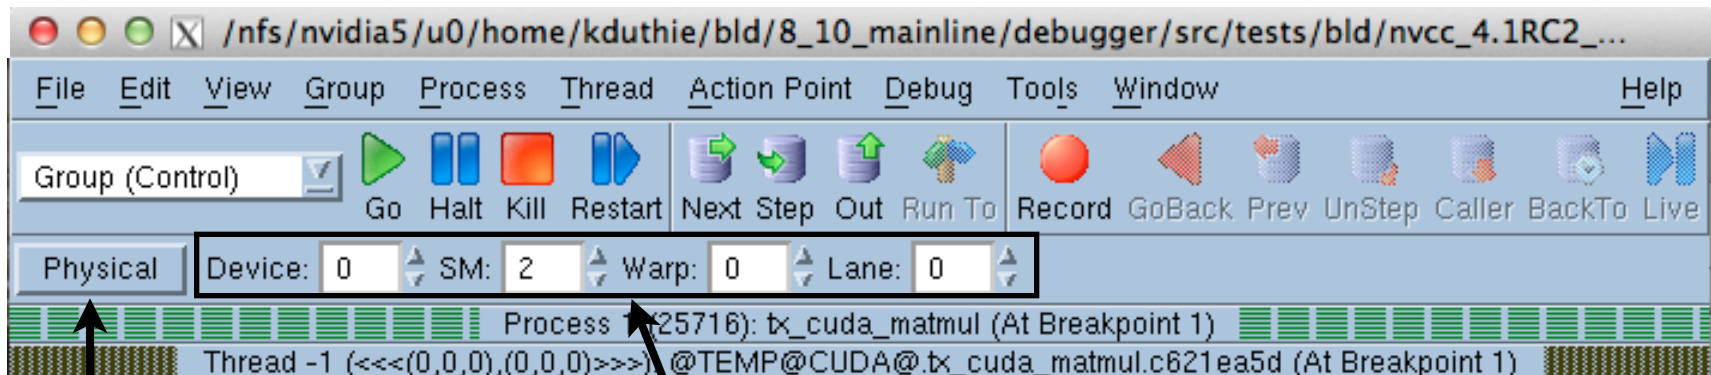
Thread
(x,y,z)

User-controlled “spinboxes” allow selection and display of any part of your GPU execution

GPU focus thread selector for changing the logical block and thread indexes of the CUDA thread.

- Logical: 2 or 3D Grid of Blocks, 3D Thread Within Grid

Debugging CUDA - Navigation



Device, SM, Warp,
and Lane

User-controlled “spinboxes” allow selection and display of any part of your GPU execution

GPU focus selector for changing physical indexes
of the CUDA thread.


- Physical: Device, SM, Warp, Lane

Executing GPU Code - Threads and Warps

- **Single-step operation advances all of the GPU hardware threads in the *same* warp**
- **To advance the execution of more than one warp:**
 - set a breakpoint and continue the process, or
 - select a line number in the source pane and select “Run To”.
- **Warps advance synchronously**
 - Warps share a PC
- **Single stepping**
 - Advances the warp containing the focus thread
 - Stepping over a `__syncthreads()` call advances all the relevant threads
- **Continue and runto**
 - Continues more than just the warp
- **Halt**
 - Stops all the host and device threads

CUDA Segmentation Faults

- **TotalView displays segmentation faults as expected**
 - **Enable CUDA memory checking in New Program dialog window**



OpenACC

The screenshot displays the OpenACC development environment running on a Windows Vista desktop. It consists of three main windows:

- Terminal Window (n15765@vista:~\$ ACC_tests/ftn_test):** Shows the output of the compilation process. It details the mapping of ELF string data, the digestion of ELF symbols, and the linking of libraries. Key messages include:
 - Mapping 981 bytes of ELF string data from '@TEMP@CUDA@_man.8df39e39'...
 - Digesting 77 ELF symbols from '@TEMP@CUDA@_man.8df39e39'...
 - Library @TEMP@CUDA@_man.49d7eedc, with 1 asects, was linked at 0x000003d5, and initially loaded at 0xFF000000a3c43100.
 - INFO: Copying library "@TEMP@CUDA@_man.49d7eedc" into the local file cache ...
 - Mapping 13 bytes of ELF string data from '@TEMP@CUDA@_man.49d7eedc'...
 - Digesting 11 ELF symbols from '@TEMP@CUDA@_man.49d7eedc'...
 - Library @TEMP@CUDA@_man.54051f77, with 5 asects, was linked at 0x0000124c, and initially loaded at 0xFF000000a3c43200.
 - INFO: Copying library "@TEMP@CUDA@_man.54051f77" into the local file cache ...
 - Mapping 167 bytes of ELF string data from '@TEMP@CUDA@_man.54051f77'...
 - Digesting 47 ELF symbols from '@TEMP@CUDA@_man.54051f77'...
 - Indexing 2040 bytes of DWARF '.debug_frame' symbols from '@TEMP@CUDA@_man.54051f77'...
 - Skimming 1611 bytes of DWARF '.debug_info' symbols from '@TEMP@CUDA@_man.54051f77'...
 - Incorporating 1600 bytes of DWARF '.debug_info' information for man.f90.i (line number)...
 - Incorporating 1600 bytes of DWARF '.debug_info' information for man.f90.i (symbol)...
- TotalView 8X.10.0-6A Window:** Displays a table of active processes.

ID	Rank	Host	Status	Description
1	<local>		T	aprun (2 active threads)
2	0 nid00130		T	aprun<man>.0 (3 active threads)
- aprun<man>.0 Window:** Shows the source code for the function `test_openacc_sck_L11_1` in `man.f90`. The code includes a parallel loop for setting variables `a`, `b`, and `c`. The stack trace indicates the function is running on Rank 0.


```

PROGRAM test_openacc
  IMPLICIT NONE
  INTEGER, PARAMETER :: M=100000
  INTEGER :: a(M), b(M), c(M)
  INTEGER :: j, total, expected


  !!$ For simple cases, use parallel loop as a shortcut for
  !!$ parallel and loop
  !!$ Set a,b,c

  !$acc parallel loop
    DO j = 1,M
      a(j) = j
      b(j) = j
      c(j) = j
    ENDDO
  !$acc end parallel loop

  !!$ Set b, copy it to host
  !$acc parallel copyout(b)
  !$acc loop
    DO j = 1,M
      b(j) = j
    ENDDO
  !$acc end loop

  !!$ Set c, copy it to host
  !$acc parallel copyout(c)
  !$acc loop
    DO j = 1,M
      c(j) = -j
    ENDDO
  !$acc end loop
  !$acc end parallel
      
```


Intel Phi Debugging with TotalView



A Spectrum of Programming Use Models

Xeon-Centric

MIC-Centric

Xeon-native

Offload

Symmetric

Reverse Offload

MIC-native

General Purpose
Serial and Parallel
Codes

Codes with
balanced needs

Highly parallel
codes

Main()
MPI_
Foo()

Scalar codes
with highly
parallel phases

Main()
Foo()

Parallel codes
with scalar
phases

Main()
MPI_
Foo()

Main()

offload<MIC>
Foo()

Main()
Foo()

Main()

Offload<Xeon>
Foo()

Intel MIC Port of TotalView

Process 2 (5856@192.168.1.100): offload_main (Mixed)
Thread 2 (139985823807232) (At Breakpoint 6)

Stack Trace

- compute07, FP=7f50fd4d24f0
- L_sample07_76__par_region1_2_39, FP=7f50fd4d24d2
- __offload_entry_sample07_c_76sample07, FP=7f50fd4d24d2
- _ZN17OffloadDescriptor7offloadEjPPv50_tS0_t, FP=7f50fd4d24d2
- _COISinkPipe::RunFunction, FP=7f50fd4d24d2
- _COISinkPipe::ProcessMessages, FP=7f50fd4d24d2
- _COISinkPipe::ThreadProc, FP=7f50fd4d24d2
- start_thread, FP=7f50fd4d24d2
- _clone, FP=7f50fd4d24d2

Stack Frame

Function "compute07":

- out: 0x7f50fd4d2754 -> 0x41400000 (109c)
- size: 0x00000010 (16)
- Local variables:
- i: 0x00000010 (16)

Registers for the frame:

- %rax: 0x7f50fd4d2754 (139985823803220)
- %rdx: 0x00000010 (16)
- %rcx: 0x7f50fd4d2754 (139985823803220)
- %rbx: 0x7f50fd4d2754 (139985823803220)

Function compute07 in sample07.c

```
90 for (i=0; i<s; i++)
91 {
92     array1[i] = p[i];
93 }
94
95 #ifdef __MIC__
96     retval = 1;
97 #else
98     retval = 0;
99 #endif
100
101 // Return 1 if array initialization was done on target
102 return retval;
103 }
104
105 __attribute__((target(mic))) void compute07(int* out, int size)
106 {
107     int i;
108     for (i=0; i<size; i++)
109     {
110         out[i] = array1[i]*2;
111     }
112 }
113 //.....
114
```

File Edit View Tools Window Help

ID	Rank	Host	Status	Description
1		<local>	R	/opt/intel/composerxe/Sample
1.1		<local>	R	in main
1.2		<local>	R	in __poll
1.3		<local>	R	in __poll
1.4		<local>	R	in pthread_cond_wait
2		192.168.1.100	M	/tmp/coi_procs/1/5856/offload
2.1		192.168.1.100	R	in sem_wait
2.2		192.168.1.100	B	in compute07
2.3		192.168.1.100	R	in __poll
2.4		192.168.1.100	R	in pthread_cond_wait

- Full visibility of both host and Phi threads
- Full support of MPI programs
- Symmetric Debugging of heterogeneous applications with offloaded code
- Remote Debugging of Phi-native applications
- Asynchronous thread control on both Xeon and Phi

Debugging MPI Applications

The screenshot displays the Rogue Wave Software's MPI debugging interface. The main window shows a C source file 'tx_basic_mpi.c' with a breakpoint set at line 106. The 'Stack Trace' shows the current frame is 'main'. The 'Process' pane shows 20 processes attached. The 'Thread' pane shows the current thread is 'Thread 1 (140313820813056): tx_basic_mpi (At Breakpoint 2)'. The 'Action Points' pane shows a list of processes from p1 to p19.

Process List:

Attach	Host	Comm	Rank	Program
<input checked="" type="checkbox"/>	192.168.1.100	0	0	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	1	1	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	2	2	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	3	3	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	4	4	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	5	5	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	6	6	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	7	7	/tmp/tx_basic_mpi
<input checked="" type="checkbox"/>	192.168.1.100	8	8	/tmp/tx_basic_mpi

Filters:

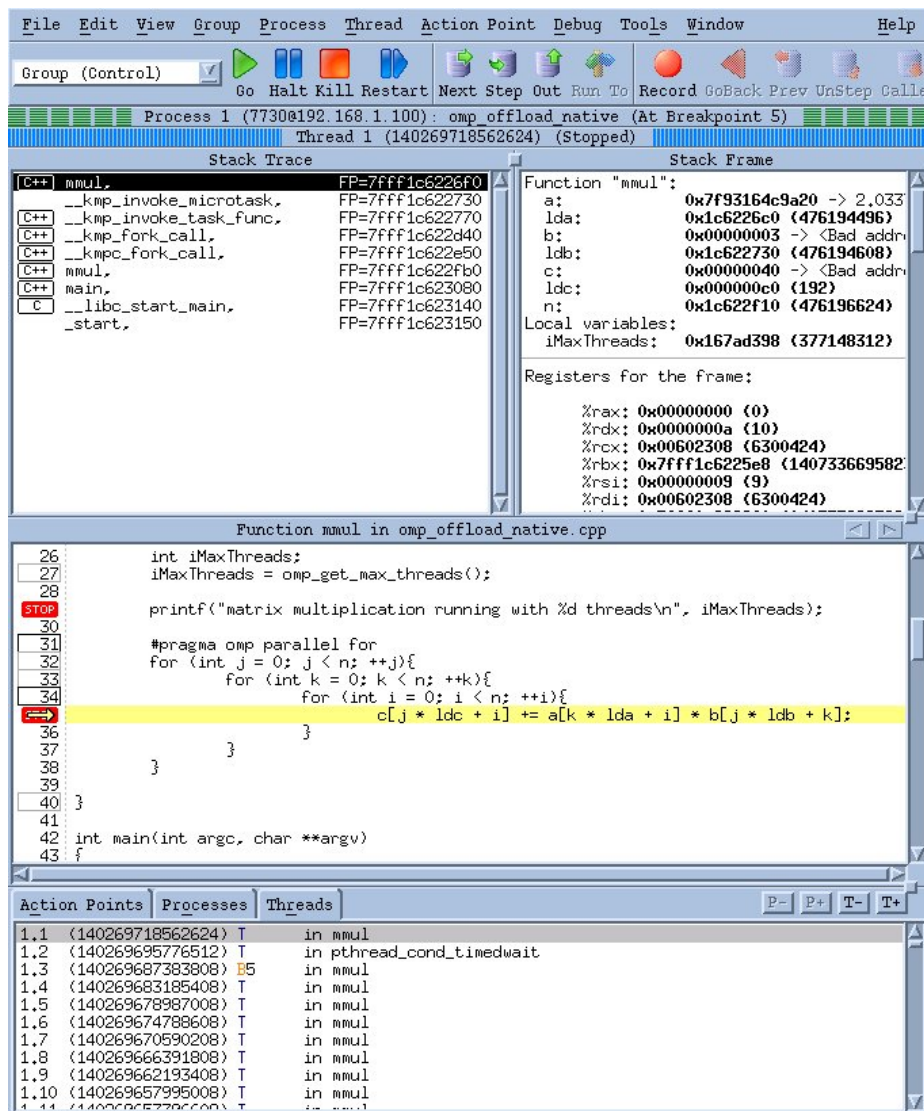
- Communicator: All
- Talking to Rank: All
- Array of Ranks: []
- List of Ranks: []
- Message Type: ☒ Send ☒ Receive ☒ Unexpected
- Halt control group: ☒

Action Points:

Action Points	Processes	Threads
p1	0	1
	2	3
	4	5
	6	7
	8	9
	10	11
	12	13
	14	15
	16	17
	18	19

- Attach to subset of processes on Phi
- Set breakpoints
- Debug MPI “as usual”

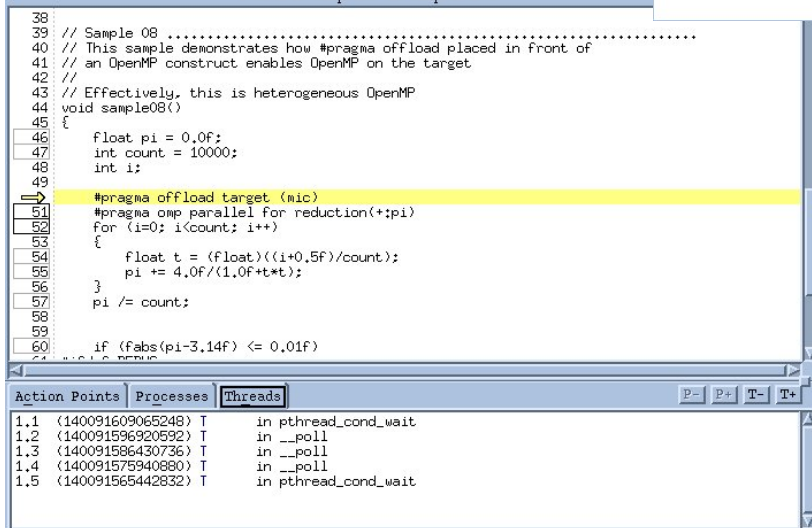
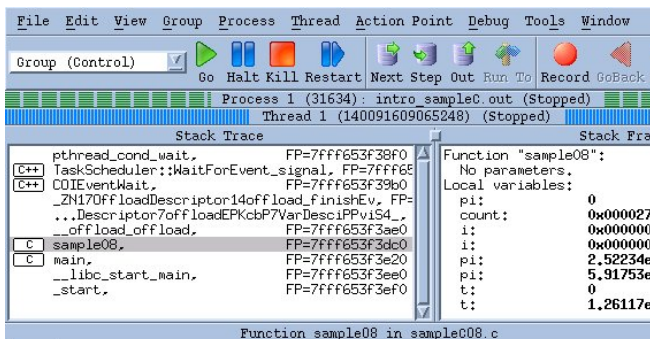
Remote Debugging of Applications on Phi



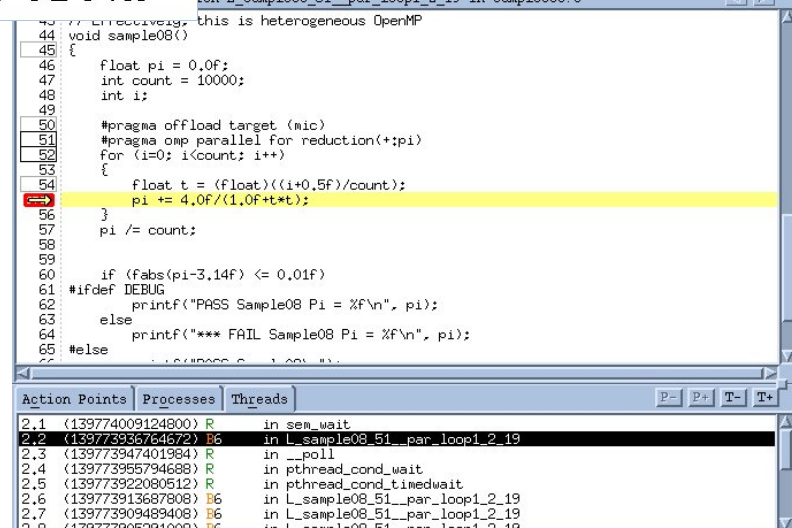
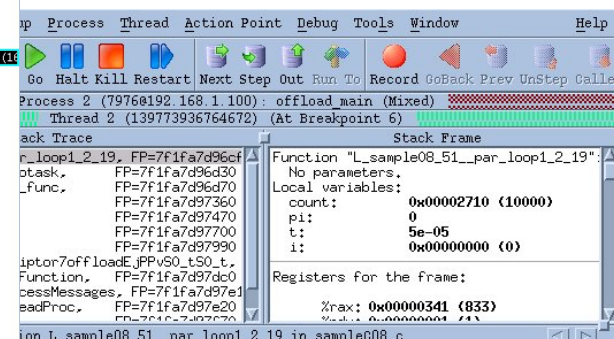
- Start application on Phi card
- Attach to running application
- See thread private data
- Investigate individual threads
- Kill stuck processes on Phi

Debugging Applications with Offloaded Code

Xeon side



Phi side



One debugging session for Phi-accelerated code



TotalView provides a full spectrum of debugging solutions

The logo for TotalView, featuring the word "TotalView" in a bold, sans-serif font with a registered trademark symbol. To the left of the text is a stylized blue wave graphic. The entire logo is enclosed within a white oval with a subtle drop shadow.

TotalView®

Code debugging

- **Highly scalable interactive GUI debugger**
 - Easy to use -- without sacrificing detail that users need to debug
 - Used from workstations to the largest supercomputers
- **Powerful features for debugging multi-threaded, multi-process, and MPI parallel programs**
- **Compatible with wide variety of compilers across several platforms and operating systems**

- **Memory Debugging**

- **Parallel memory analysis and error detection**
 - Intuitive for both intensive and infrequent users
- **Easily integrated into the validation process**

- **Reverse Debugging**

- **Parallel record and deterministic replay within TotalView**
 - Users can run their program “backwards” to find bugs
- **Allows straightforward resolution of otherwise stochastic bugs**

- **GPU CUDA Debugging**

- **Full Hybrid Architecture Support**
- **Asynchronous Warp Control**
- **Multi-Device and MPI Support**

Thanks!



Thank You

Download a TotalView evaluation at:
[*www.roguewave.com/products*](http://www.roguewave.com/products)
[*ehinkel@roguewave.com*](mailto:ehinkel@roguewave.com)